



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

CERTIFIED COPY OF PRIORITY DOCUMENT



the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name under which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., p.l.c., L.C. or PLC.


Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 22 June 2001

THIS PAGE BLANK (USPTO)

Patents Form 1/77

Patents Act 1977
(Rule 16)The
Patent
Office05JUL00 E550380-1 000393
P01/7700 0.00-0016476.4

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

2000P04865/GB/R76/MM/rr

2. Patent application number

(The Patent Office will fill in this part)

0016476.4

05 JUL 2000

3. Full name, address and postcode of the or of each applicant (underline all surnames)

ROKE MANOR RESEARCH LIMITED
OLD SALISBURY LANE
ROMSEY
SO51 0ZN
UNITED KINGDOM

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

5615455006
UNITED KINGDOM

4. Title of the invention

TRANSMITTING CONTROL INFORMATION ACROSS
A ROUTING DEVICE

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom
to which all correspondence should be sent
(Including the postcode)

DEREK ALLEN

Siemens Shared Services Limited
Intellectual Property Department
Siemens House, Oldbury
Bracknell, Berkshire RG12 8FZ
United Kingdom

Patents ADP number (if you know it)

7396419002

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

NO

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))


Patents Form 1/77

THIS PAGE BLANK (USPTO)

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form 0
Description 7 ✓
Claim(s) 0
Abstract 0
Drawing(s) 2 ✓



10. If you are also filing any of the following, state how many against each item.

Priority documents


Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application.
Signature  Date 05.07.2000
Margaret MACKETT

12. Name and daytime telephone number of person to contact in the United Kingdom
Margaret Mackett - 01344 396808

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

Patents Form 1/77

THIS PAGE BLANK (USPTO)

TRANSMITTING CONTROL INFORMATION ACROSS A ROUTING DEVICE

The present invention relates to a technique for transmitting control
5 information across a routing device. For the purposes of the following
description the term routing device refers to any device which performs the
function of a circuit switch or a router.

Data is transferred over the Internet by means of a plurality of
routing devices in accordance with a standard protocol known as Internet
10 Protocol (IP). IP is a protocol based on the transfer of data variable sized
portions known as packets. All network traffic involves the transportation
of packets of data. Routers are devices for accepting incoming packets;
temporarily storing each packet; and then forwarding the packets to another
part of the network.

15 Traffic volume in the Internet is growing exponentially, almost
doubling every 3 months, and the capacity of conventional IP routers is
insufficient to meet this demand. There is thus an urgent requirement for
products that can route IP traffic at extremely large aggregate bandwidths
in the order of several terabits per second. Such routing devices are termed
20 "terabit routers".

Terabit routers require a scalable high capacity communications path
between the point at which packets arrive at the router (the "ingress") and
the point at which the packets leave the router (the "egress"). The terabit
routers discussed hereafter are intrinsically optical routing devices, since
25 light is an effective and efficient medium for such high capacity
communications.

The packets transferred in accordance with IP can (and do) vary in size. Within routers it has been found useful to pass data in fixed sized units. In routers then data packets are partitioned into small fixed sized units, known as cells. In discussion of routers, cells are often subdivided
5 into units of slices (64 bits) or octets (8 bits). A typical cell size is 64 octets, however cells of 128 octets are also known.

One suitable technique for implementing a scalable communications path is an optical backplane device, known as a cell based cross-bar. Data packets are partitioned into cells by a plurality of ingress line function
10 means for passage across the cross-bar.

The plurality of ingress line function means provide respective interfaces between incoming communications channels carrying incoming data and the optical backplane. Similarly a plurality of egress line function means provide respective interfaces between the optical backplane and
15 outgoing communications channels carrying outgoing data.

A general terabit router architecture bears some similarity to conventional router architecture. Packets of data arrive at input port(s) of ingress line function means and are routed as cells across the cross-bar to a predetermined egress line function means which reassembles the packets
20 and transmits them across its output port(s). Each ingress line function means maintains a separate packet queue for each egress line function means.

Line function means are implemented as line interface cards (LICs). Since one of the line functions regularly undertaken by line function means
25 is forwarding, LICs are also known as 'forwarders'. Further line functions include congestion control and maintenance of external interfaces, input ports and output ports.

The backplane has the following features. Firstly, data is transmitted across the backplane in slices. Ingress and egress line function means are synchronised so that they each send or received slices simultaneously. At each slice time, each ingress line function means will transmit a slice which
5 can be received by one or more egresses line function means. Likewise, at each slice time, each egress line function means can receive a slice from one and only one ingress line function means. The egress line function means is responsible for selecting the correct slice.

User data is conveyed across the backplane as cells consisting of
10 some fixed integral number of slices.

The co-ordination of the transmission and reception of cells is performed by a cross-bar controller.

At the cell level, the combination of backplane and cross-bar controller functions as a cell-based cross-bar.

15 In a conventional cell based cross-bar each ingress line function means is connected to one or more of the egress line function means. However, each ingress line function means is only capable of connecting to one egress line function means at any one time. Likewise, each egress line function means is only capable of connecting to one ingress line function
20 means at a time.

All ingress line function means transmit in parallel and independently across the cross-bar. Furthermore cell transmission is synchronised with a cell cycle, having a period of, for example, 108.8ns.

The ingress line function means simultaneously each transmit a new
25 cell with each new cell cycle.

The pattern of transmissions from the ingress line function means across the cross-bar to the egress line function means changes at the end of every cell cycle.

A cross-bar controller is provided for efficient allocation of the bandwidth across the cross-bar. It calculates the rates that each ingress line function means must transmit to each egress line function means. This is the same as the rate at which data must be transmitted from each packet queue. The calculation makes use of real-time information, including traffic measurements and indications from the ingress line function means. The indications from the ingress line function means include monitoring the current rates, queue lengths and buffer full flags. The details of the calculation are discussed more rigorously in the copending UK Patent Application Number 9907313.2 (docket number F21558/98P4863).

The cross-bar controller performs a further task; it serves to schedule the transfer of data efficiently across the cross-bar whilst maintaining the calculated rates. At the end of each cell cycle, the cross-bar controller communicates with the ingress and egress line function means as follows. Firstly, the cross-bar controller calculates and transmits to each ingress line function means the identity of the next packet queue from which to transmit. Secondly, the cross-bar controller calculates and transmits to each egress line function means the identity of the ingress from which it must receive.

The architecture described above gives rise to two requirements.

The need for a means for each ingress forwarder to transmit traffic measurements and indications to the cross-bar controller and the need for a means for the cross-bar controller to send configuration information to each ingress and each egress forwarder.

It is possible to provide dedicated communications paths to meet these requirements. However such a solution requires additional hardware, which is expensive in terms of increased power consumption, installation and materials.

It is therefore an object of the invention to obviate or at least mitigate the aforementioned problems.

In accordance with the present invention, there is provided a router device having a plurality of ingress line function means, a plurality of egress line function means, a backplane and a controller means, wherein the transmission of signals from the plurality of ingress line function means to the controller means and signals from the controller means to each of the ingress line function means and each of the egress line function means takes place across the backplane.

For a better understanding of the present invention, reference will now be made, by way of example only, to the accompanying drawings in which:-

Figure 1 illustrates a terabit router architecture; and
Figure 2 shows a cross-bar controller.

A conventional terabit router architecture is depicted in Figure 1. Packets of data arrive at ingress line function means (ingress forwarders) via their input port(s) and are routed across the cross-bar to the correct egress line function means (egress forwarders) which transmits them across its output port(s). Each ingress line function means maintains a separate packet queue for each egress line function means, for example q_{11} , q_{12} and q_{13} .

A conventional cell based cross-bar is shown in Figure 2. Here each ingress forwarder may be connected to one or more of the egress forwarders. However each ingress forwarder may only be connected to one egress forwarder at a time and each egress forwarder may only be connected to one ingress forwarder at a time.

The co-ordination of the transmission and reception of cells is performed by a cross-bar controller.

Each ingress forwarder communicates traffic measurements and notifications for the use of the cross-bar controller. The cross-bar controller allocates connections between ingress and egress forwarders and informs the respective forwarders accordingly for each cell cycle in turn.

5 For transmission of control information from ingress to the controller, each ingress line function means is assigned a dedicated slice timeslot which it uses to send information to the controller. The timeslots do not overlap. When the timeslot assigned to a given ingress line function means is reached, the given ingress line function means transmits a slice of control information, interrupting its transmission of user data. The controller means selects the ingress line function means from which to receive control information according to the current timeslot number.

10 When receiving user data from a given ingress line function means, an egress line function means ignores information in a slice timeslot, if that timeslot is assigned to the given ingress line function means for transmission of control information. The position of the control slice timeslot is determined by fixed global information, for example the position of the line function means in a physical rack of LICs. This makes it simple for each line function means to determine which slice timeslot is used by each line function means for this purpose.

20 For transmission of control information from the controller to ingress and egress line function means, the same technique is used except that each line function means is assigned a dedicated timeslot on which to receive.

25 It will be readily understood that although the preceding discussion has been in terms of optical terabit routers, the apparatus of the present invention are capable of implementation in a wide variety of routing devices, including switches and routers, and that these routing devices can be either purely electronic, part electronic/part optical or optical in nature.

Further considerations when implementing the present invention, include the need for a degree of redundancy. To prevent a potentially catastrophic failure of a routing device, an further auxiliary management card having identical features to the management card can be provided to
5 replace the management card if it were to fail. Since the operation of the routing device of the present invention is essentially stateless, the auxiliary management card can replace the management card without substantial interruption.

THIS PAGE BLANK (USPTO)

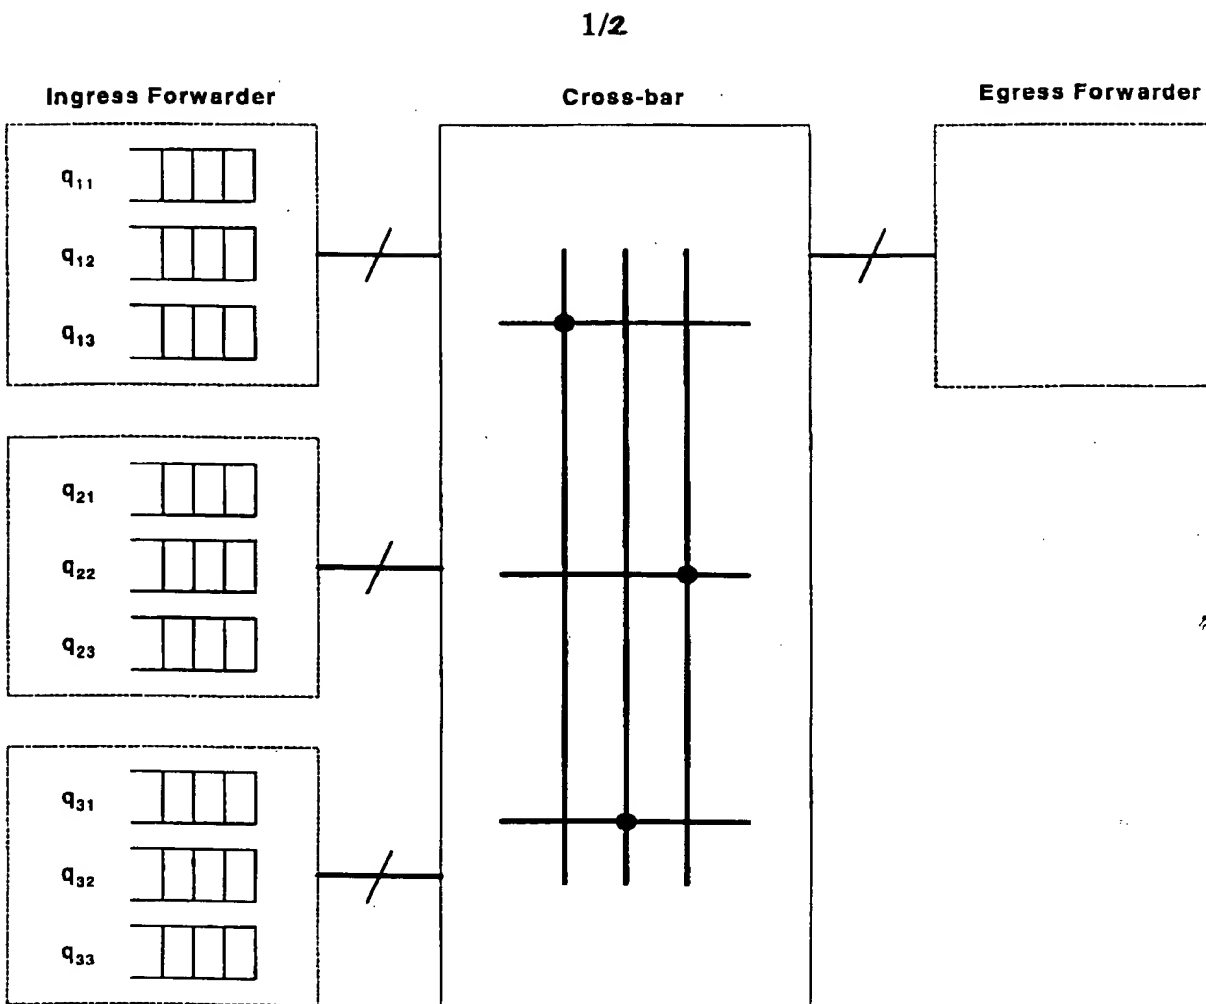
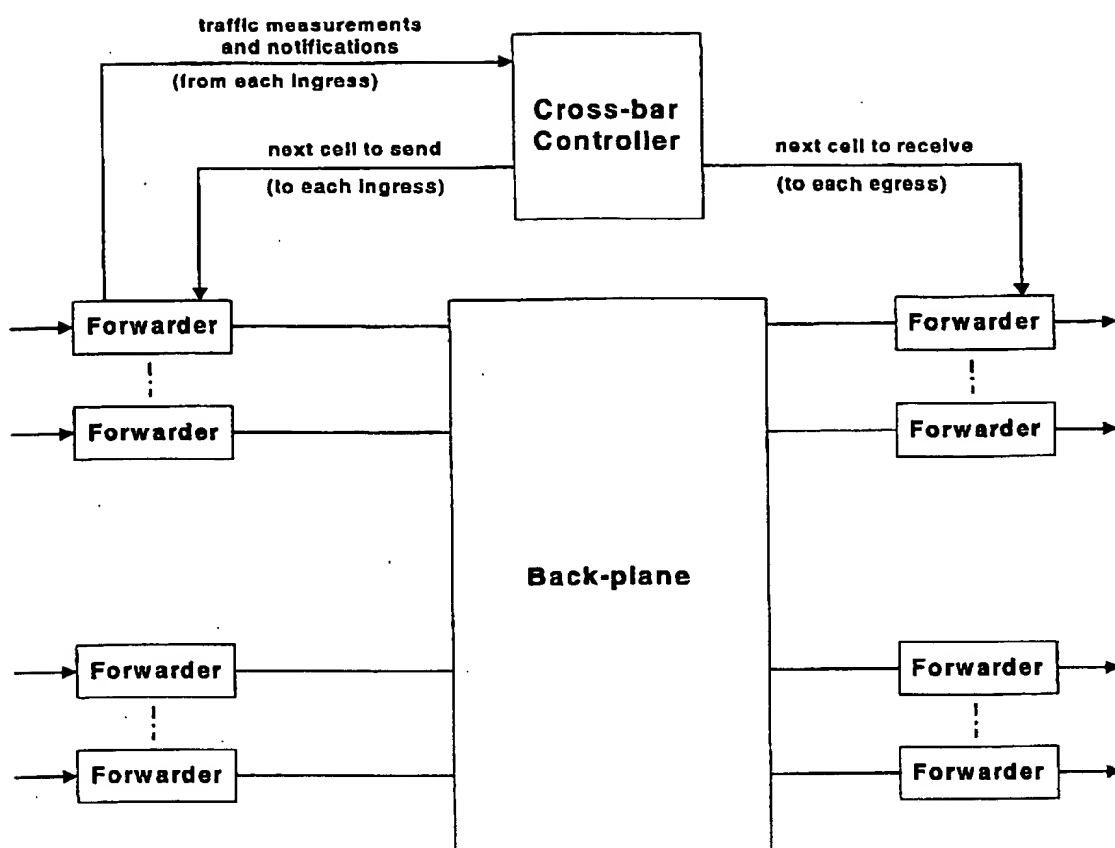


Figure 1: Terabit Router Architecture

THIS PAGE BLANK (USPTO)

2/2

**Figure 2: Cross-bar Controller**

THIS PAGE BLANK (USPTO)

THIS PAGE BLANK (USPTO)